

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES


Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.



(19)

Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 0 841 626 A1

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
13.05.1998 Bulletin 1998/20

(51) Int Cl.⁶: G06F 17/30, G06F 9/46

(21) Application number: 97308501.2

(22) Date of filing: 24.10.1997

(84) Designated Contracting States:
**AT BE CH DE DK ES FI FR GB GR IE IT LI LU MC
NL PT SE**
Designated Extension States:
AL LT LV RO SI

(72) Inventors:
• **Sealby, Arthur George**
Sandback, Cheshire, CW11 1YZ (GB)
• **Griffiths, Lawrence Gregory**
West Didsbury, Manchester, M20 2WL (GB)

(30) Priority: 12.11.1996 GB 9623472

(74) Representative: **Guyatt, Derek Charles et al**
Intellectual Property Department
International Computers Limited
Cavendish Road
Stevenage, Herts, SG1 2DY (GB)

(71) Applicant: **INTERNATIONAL COMPUTERS
LIMITED**
Putney, London, SW15 1SW (GB)

(54) Data transfer mechanism

(57) A mechanism is described for transferring data between a pair of relational databases on client and server machines. A client component creates a first pipe in the client, and sends information to a server component, enabling it to create a second pipe in the server. The client component then calls the data unloading utility of the client database, causing it to start to unload the data into the first pipe. The client component then

starts up a Send process in the client, which reads data from the first pipe and sends it to the server. At the server, the server component calls the loading utility of the server database, specifying the second pipe as input. As the data arrives at the server, the server component writes it into the second pipe, from where it will be read by the loading utility and written into the server database.

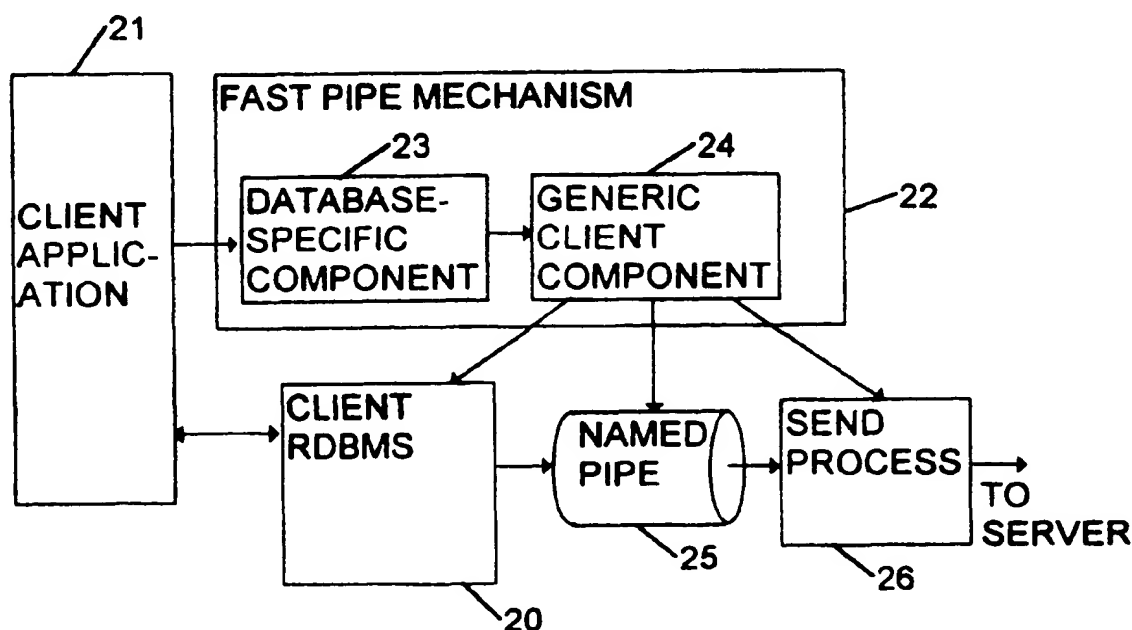


FIG. 2

EP 0 841 626 A1

Description

Background to the Invention

5 This invention relates to a data transfer mechanism. The invention is particularly, although not exclusively, concerned with a data transfer mechanism for transferring data between relational databases on client and server machines.

Conventional techniques for transferring data between two relational databases require the data first to be read from the source database into a flat file and then to be transferred from the flat file to the target database. However, a
10 problem with this is that this process tends to be slow, and so can take an excessively long time if large amounts of data are involved.

The object of the present invention is to provide an improved data transfer mechanism.

Summary of the Invention

15 According to the invention a computer system comprises a first computer containing a first database, a second computer containing a second database, and a data transfer mechanism for transferring data from the first database to the second database, wherein the data transfer mechanism comprises:

- 20 (a) in said first computer, a first pipe, an unloading process for reading data from the first database into the first pipe, and a send process for reading data from the first pipe and sending it to the second computer; and
- (b) in said second computer, a second pipe, a receiving process for writing data received from the first computer into the second pipe, and a loading process for writing data from the second pipe into the second database.

Brief Description of the Drawings

Figure 1 is a schematic block diagram of a computer system including client and server machines.

Figure 2 is a schematic block diagram showing software components running in the client machine.

Figure 3 is a schematic block diagram showing software components running in the server machine.

30 Figure 4 is a flow chart showing the operation of a generic client component.

Figure 5 is a flow chart showing the operation of a generic server component.

Description of an Embodiment of the Invention

35 One embodiment of the invention will now be described by way of example with reference to the accompanying drawings.

Referring to Figure 1, a computer system comprises a client machine 10 and a server machine 11, connected by a network 12. In the present example, the client and server machines are both UNIX machines. For example, the server machine may be an ICL Goldrush machine. The network may use the standard TCP/IP protocol.

40 UNIX allows the use of named pipes. A named pipe is a standard UNIX mechanism for redirecting the output of one process to the input of a second process. Pipes are treated as normal files by UNIX: they can be accessed by code using standard IO system calls.

Referring to Figure 2, the client machine includes a relational database management system (RDBMS) 20, and a client application 21. The RDBMS may, for example, comprise an Ingres or Oracle database. The client application
45 may be any application which makes use of the RDBMS.

Referring to Figure 3, the server machine also includes an RDBMS 30. The server RDBMS 30 may be of the same kind as the client RDBMS 20, or may be different. For example, both may be Ingres databases, or the client RDBMS may be an Ingres database and the server RDBMS may be an Oracle database. It is assumed that it is required to
50 transfer data from the client RDBMS 20 to the server RDBMS 30.

Referring again to Figure 2, the client machine includes a fast pipe mechanism 22 for efficient transfer of data from the client RDBMS to the server RDBMS. The fast pipe mechanism consists of two components: an database-specific component 23 and a generic client component 24. The database-specific component 23 is specific to the particular
55 client RDBMS, and the client may include several different database-specific components for different RDBMS types. For example, one specific component may be provided for Ingres and another for Oracle. The generic client component 24, on the other hand, is generic to all RDBMS types.

In operation, when the client application 21 requires to transfer data from the client RDBMS to the server RDBMS,

it makes a call to the database-specific component 23. The call interface to this component has parameters which are specific to the unloading utilities offered by the client RDBMS and to the loading utilities offered by the server RDBMS. By way of example, the database-specific component specific to Ingres databases provides a command **"ingtooracle"** for transferring a table from an Ingres database to an Oracle database. This command has the following parameters:

5

-ppathname	The pathname of the server (i.e. the network address of the server machine).
-srows	(Optional) The number of rows from the beginning of the table not to be loaded into the server database. This is useful if it is required to continue a load which has been interrupted for some reason.
-ddirectory	(Optional) The name of a directory which holds "copydb" script files.
userid	The username/password of the user performing the load on the server machine.
dbname	The name of the database on the client machine which holds the table to be transferred.
tablename	The name of the table in the specified database which is to be transferred.

10

15

The database-specific component validates the parameters of the call, and then constructs a **"fastpipe"** call to the generic client component 24, containing the following parameters:

20

-ppathname	The pathname of the server.
-hheader	A header (see below)
-fdata_filename	This specifies the name of a pipe to be used for transferring data from the client database.
-ssql_filename	The name of a file containing an SQL script which may be required by the loading utility in the server.
client_command_line	A command line, for execution by the client.

25

The header is in "tld" (type, length, data) format, and includes the following information:

30

- A value specifying whether the data to be transferred is formatted or unformatted.
- The SQL script filename (if one has been specified). This duplicates the information in the parameter **-ssql_filename**. A length of 0 indicates that no SQL script is supplied.
- A server command line, for execution by the server.
- A pipe name for use by the server.

35

When the generic client component receives this call, it performs the following actions as shown in Figure 4:

(Step 41) The generic client component creates a named pipe 25, with the name **-fdata_filename**.

(Step 42) The generic client component then sends the header over the network 12 to the server, using the specified

-ppathname as the network destination address.

(Step 43) If an SQL script filename has been specified, the generic client component then accesses this script and sends it to the server, immediately after the header.

(Step 44) The generic client component then executes the client command line. This calls the data unloading utility of the client RDBMS, specifying the named pipe 25 as the output file. The unloading utility then starts to unload the

40

required data from the client database and to write it into the named pipe 25.

(Step 45) It starts up a send process 26. This reads the data from the named pipe 25, copies it into a buffer and, when the buffer is full, sends the data over the network to the server.

Referring again to Figure 3, the server machine includes a generic server component 31. This receives the header sent by the client, as well as any SQL script, and performs the following actions, as shown in Figure 5:

45

(Step 51) The generic server component extracts the pipe name from the header, and uses this to create a named pipe 32.

(Step 52) The generic server component then extracts the SQL script filename (if any) from the header, and uses this to load and run the specified SQL script.

(Step 53) The generic server component then extracts the server command line extracted from the header, and executes it. This calls the data loading utility in the server RDBMS, specifying the named pipe 33 as the input file. In the case where the server RDBMS is an Oracle database, the server command line contains the **userid** parameter as a password to allow it to access the required tables.

50

(Step 51) As the data arrives from the network, the generic server component writes this data into the pipe 32. The loading utility will then read this data from the pipe 32 and write it into the database.

55

It can be seen that the use of the pipes 25, 32 in this way allows the send process 26 to start transferring data over the network as soon as some data is available in the pipe 25, without having to wait for the whole table to be unloaded. Also, the loading process on the remote server machine begins as soon as the pipe 32 receives the first packet of data from the network. Hence, the unloading, transfer and loading processes all proceed simultaneously (or

in parallel). This greatly improves the efficiency of the transfer mechanism.

Some possible modifications

It will be appreciated that many modifications may be made to the system described above without departing from the scope of the present invention. For example, compression and decompression techniques may be used in transmitting the data.

Claims

1. A computer system comprising a first computer (10) containing a first database (20), a second computer (11) containing a second database (30), and a data transfer mechanism for transferring data from the first database to the second database, characterised in that the data transfer mechanism comprises:

(a) in said first computer, a first pipe (25), an unloading process for reading data from the first database into the first pipe, and a send process (26) for reading data from the first pipe and sending it to the second computer; and

(b) in said second computer, a second pipe (32), a receiving process (31) for writing data received from the first computer into the second pipe, and a loading process for writing data from the second pipe into the second database.

2. A system according to Claim 1 wherein the data transfer mechanism includes:

(a) in said first computer, means (23) for creating first and second command lines, means (44) for causing the first command line to be executed by the first computer, to initiate the unloading process, and means (42) for sending the second command line to the second computer; and

(b) in said second computer, means (53) for executing the second command line to initiate said loading process.

3. A data processing system comprising a first computer (10) including a first database (20) having a data unloading utility, a second computer (11) including a second database (30) having a data loading utility, and a data transfer mechanism for transferring data from the first database to the second database, characterised in that the data transfer mechanism comprises:

(a) in the first computer, means (41) for creating a first pipe (25), means (42) for sending a message to the second computer, specifying a second pipe (32) in said second computer, means (44) for calling said data unloading utility of said first database, specifying the first pipe as an output, and means (45) for starting up a send process (26) in said first computer, to read data from said first pipe and send the data to said second computer; and

(b) in said second computer, means (51) responsive to said message, for creating said second pipe (32) in said second computer, means (53) for calling said data loading utility of said second database, specifying the second pipe as an input, and means (54) for receiving said data from said send process and writing said data into said second pipe.

4. A system according to Claim 3 wherein the data transfer mechanism includes:

(a) in said first computer, means (23) for creating first and second command lines, means (44) for causing the first command line to be executed by the first computer, to call said data unloading utility, and means (42) for sending the second command line to the second computer; and

(b) in the second computer, means (53) for executing the second command line to call said data loading utility.

5. A method of transferring data between a first database on a first computer and a second database on a second computer, characterised in that the method comprises the steps:

- (a) creating a first pipe in said first computer;
(b) sending a message to the second computer, specifying a second pipe to be created in said second computer;
(c) creating said second pipe in said second computer;
(d) in said first computer, unloading data from said first database into said first pipe;
(e) in said first computer, reading data from said first pipe and sending the data to said second computer;
(f) receiving said data in said second computer, and writing said data into said second pipe; and
(g) in said second computer, reading data from said second pipe and loading it into the second database.

6. A method according to Claim 5, including the further steps:

- (a) creating first and second command lines in the first computer;
(b) causing the first command line to be executed by the first computer, to initiate unloading of data from said first database into said first pipe;
(c) sending the second command line to the second computer; and
(d) executing the second command line in the second computer, to initiate loading of data into the second database.




FIG. 1




FIG. 2




FIG. 3

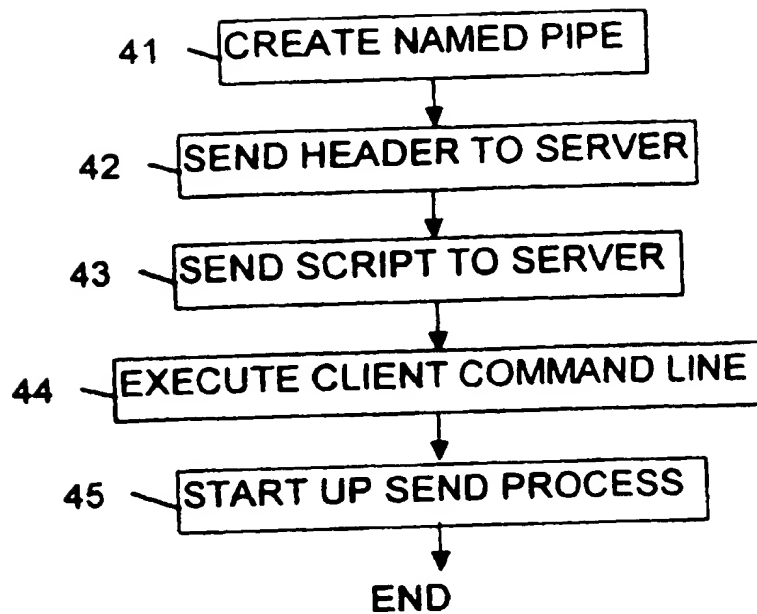


FIG. 4

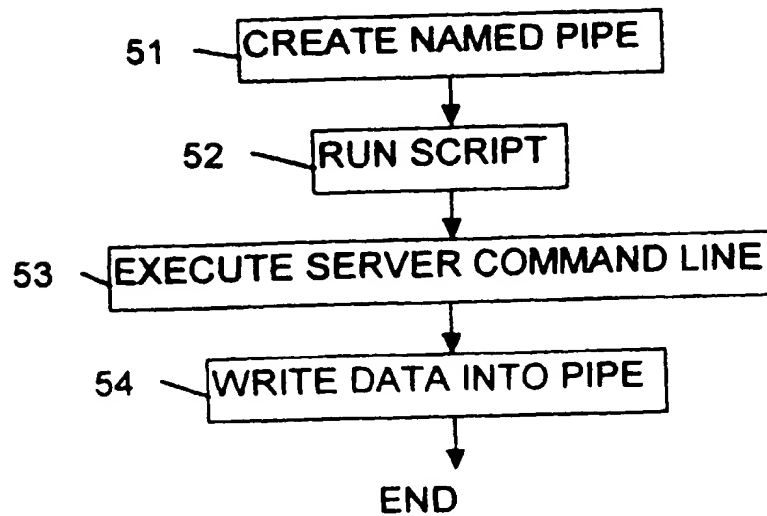


FIG. 5



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 97 30 8501

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
Y	EP 0 586 157 A (AMERICAN TELEPHONE & TELEGRAPH)	1,3,5	G06F17/30 G06F9/46
A	* the whole document *	2,4,6	
Y	DICK CONKLIN / GENERAL EDITOR: "OS/2 Notebook - The Best of IBM Personal Systems Developer" 1990, MICROSOFT PRESS, USA XP002056492	1,3,5	
A	* page 348, line 1 - page 358, last line *	2,4,6	
A	EP 0 398 641 A (IBM) * the whole document *	1-6	
			TECHNICAL FIELDS SEARCHED (Int.Cl.6)
			G06F
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 20 February 1998	Examiner Fonderson, A
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone ✓ : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons A : member of the same patent family, corresponding document</p>			

FPD FORM 1503 33 82 (P01.01)